PERM: A MONTE CARLO STRATEGY FOR SIMULATING POLYMERS AND OTHER THINGS

P. GRASSBERGER[1,2], H. FRAUENKRON[1], and W. NADLER[1]
[1] *HLRZ c/o Forschungszentrum Jülich, D-52425 Jülich, Germany*
[2] *Physics Department, University of Wuppertal, D-42097 Wuppertal, Germany*

We describe a general strategy, PERM (*P*runed-*E*nriched *R*osenbluth *M*ethod), for sampling configurations from a given Gibbs-Boltzmann distribution. The method is *not* based on the Metropolis concept of establishing a Markov process whose stationary state is the wanted distribution. Instead, it starts off building instances according to a biased distribution, but corrects for this by cloning "good" and killing "bad" configurations. In doing so, it uses the fact that nontrivial problems in statistical physics are high dimensional. Therefore, instances are built step by step, and the final "success" of an instance can be guessed at an early stage. Using weighted samples, this is done so that the final distribution is strictly unbiased. In contrast to evolutionary algorithms, the cloning/killing is done without simultaneously keeping a large population in computer memory. We apply this in large scale simulations of homopolymers near the theta and unmixing critical points. In addition we sketch other applications, notably to polymers in confined geometries and to randomly branched polymers. For theta polymers we confirm the very strong logarithmic corrections found in previous work. For critical unmixing we essentially confirm the Flory-Huggins mean field theory and the logarithmic corrections to it computed by Duplantier. We suggest that the latter are responsible for some apparent violations of mean field behavior. This concerns in particular the exponent for the chain length dependence of the critical density which is 1/2 in Flory-Huggins theory, but is claimed to be ≈ 0.38 in several experiments.

1 Introduction

For many statistical physicists, "Monte Carlo" is synonymous for the Metropolis strategy[1] where one sets up an ergodic Markov process which has the desired Gibbs-Boltzmann distribution as its unique asymptotic state. There exist numerous refinements concerned with more efficient transitions in the Markov process (e.g. cluster flips[2] or pivot moves[3]), or with distributions biased such that false minima are less harmful and that autocorrelations are reduced (e.g. multicanonical[4] sampling and simulated tempering[5]). But most of these schemes remain entirely within the framework of the Metropolis strategy.

On the other hand, there exists a very extensive literature on simulations of polymer systems by completely different methods, starting already in the mid-50's and continuing until today. Polymers represent a challenge for simulations because of the constraints imposed by chain connectivity on the one hand, and steric hindrance on the other. These can lead to entanglement effects which seriously reduce the mobility of the monomers, and which can render Metropolis type schemes inefficient.

The first such alternative method was invented by Rosenbluth and Rosenbluth.[6] They observed that the exponential attrition in straightforward simulations of unbiased self avoiding random walks (SAW's) can be strongly reduced by simulating a biased sample, and correcting the bias by means of a weight associated to each configuration. The biased sample is simply obtained by replacing any "illegal" step

which would violate the self avoidance constraint by a random "legal" one, provided such a legal step exists. More generally, assume we want to simulate a distribution in which each configuration $\mathcal{C}$ is weighted by a (Boltzmann-)weight $Q(\mathcal{C})$, so that for any observable A one has $\langle A\rangle = \sum_{\mathcal{C}} A(\mathcal{C})Q(\mathcal{C})/\sum_{\mathcal{C}} Q(\mathcal{C})$. If we sample unevenly with probability $p(\mathcal{C})$, then we must compensate this by giving a weight $\propto Q(\mathcal{C})/p(\mathcal{C})$,

$$\langle A\rangle = \lim_{M\to\infty} \frac{\sum_{i=1}^{M} A(\mathcal{C}_i)Q(\mathcal{C}_i)/p(\mathcal{C}_i)}{\sum_{i=1}^{M} Q(\mathcal{C}_i)/p(\mathcal{C}_i)} \equiv \lim_{M\to\infty} \frac{1}{M}\sum_{i=1}^{M} A(\mathcal{C}_i)W(\mathcal{C}_i)\ . \tag{1}$$

We call this the generalized Rosenbluth method. If $p(\mathcal{C})$ were chosen close to $Q(\mathcal{C})$, this would lead to importance sampling and obviously would be very efficient. But in general this is not possible, and Eq.(1) suffers from the problem that the sum is dominated by very few events with high weight.

Consider now a lattice chain of length $N+1$ with self avoidance and with nearest neighbor interaction $-\epsilon$ between unbonded neighbors. In the original Rosenbluth method, $p(\mathcal{C})$ is then a product,

$$p(\mathcal{C}) \propto \prod_{n=1}^{N} \frac{1}{m_n}, \tag{2}$$

where m_n is the number of free neighbors in the n-th step, i.e. the number of possible lattice sites where to place the n-th monomer (monomers are labeled $n = 0, 1, \ldots N$). Similarly, $Q(\mathcal{C})$ is a product,

$$Q(\mathcal{C}) = \prod_{n=1}^{N} e^{-\beta E_n}, \tag{3}$$

where $\beta = 1/kT$ and $E_n = -\epsilon\sum_{k=0}^{n-1}\Delta_{kn}$ is the energy of the n-th monomer in the field of all previous ones ($\Delta_{kn} = 1$ if and only if monomers k and n are neighbors and non-bonded, otherwise $\Delta_{kn} = 0$).

Obviously, $p(\mathcal{C})$ favors compact configurations where monomers have only few free neighbors. This renders the Rosenbluth method unsuitable for long chains, except near the collapse ('theta') point where simulations with $N \leq 1000$ are feasible on the simple cubic lattice.[7] In general we should find ways to modify the sampling so that "good" configurations are sampled more frequently, and "bad" ones less. The key to this is the product structure of the weights

$$W(\mathcal{C}_i) = \frac{Q(\mathcal{C}_i)/p(\mathcal{C}_i)}{M^{-1}\sum_{k=1}^{M} Q(\mathcal{C}_k)/p(\mathcal{C}_k)} = \frac{1}{\hat{Z}_N}\prod_{n=1}^{N} w_n(\mathcal{C}_i) \tag{4}$$

implied by Eqs. (2) and (3). Here, $\hat{Z}_N = M^{-1}\sum_{k=1}^{M} Q(\mathcal{C}_k)/p(\mathcal{C}_k)$ is an estimate of the partition sum. A similar product structure holds in practically all interesting cases.

We can thus watch how the weight builds up while the chain is constructed step by step. If the partial weight (from now on we drop the dependence on $\mathcal{C}$)

$$W_n = \hat{Z}_n^{-1} \prod_{j=1}^{n} w_j \ , \qquad 1 < n < N \ , \tag{5}$$

gets too large (i.e. is above some threshold W^+), we replace the configuration by k copies, each with weight W_n/k. One of these copies is continued to grow, all others are placed on a stack for later use. Following Ref.[8] we call this 'enrichment'. The opposite action, when W_n falls below another threshold W^- ('pruning'), is done stochastically: With probability 1/2 the configuration is killed and replaced by the top of the stack, while its weight is doubled in the other half of cases.

In this way PERM (pruned-enriched Rosenbluth method[9]) gives a sample with exactly the right statistical weights, independently of the thresholds $W^\pm$, the selection probability $p(\mathcal{C})$, and the clone multiplicity k. But its efficiency depends strongly on good choices for these parameters. Notice that one has complete freedom in choosing them, and can even change them during a run. Fortunately, reasonably good choices are easy to find (more sophisticated choices needed at very low temperatures are discussed in Ref.[10,11]). The guiding principle for $p(\mathcal{C})$ is that it should lead as closely as possible to the correct final distribution, so that pruning and enrichment are kept to a minimum. This is also part of the guiding principles for $W^\pm$. In addition, W^+ and W^- have to be chosen such that roughly the same effort is spent on simulating any part of the configuration. For polymers this means that the sample size should neither grow nor decrease with chain length n. This is easily done by adjusting W^+ and W^- 'on the fly', see Ref.[9] where a pseudocode is given for the entire algorithm.

In selecting the good and killing the bad, PERM is similar to evolutionary and genetic algorithms,[12] to population based growth algorithms for chain polymers,[13,14,15,16] to diffusion type quantum Monte Carlo algorithms,[17] and to the 'go with the winners' strategy of Ref.[18]. The main difference with the first three groups of methods is that we do not keep the entire population of instances simultaneously in computer memory. Indeed, even on the stack we do not keep copies of good configurations but only the steps involved in constructing the configurations and flags telling us when to make a copy.[9] In genetic algorithms, keeping the entire population in memory is needed for cross-overs, and it allows a one-to-one competition between instances. But in our case this is not needed since every instance can be compared to the average behavior of all others. The same would be true for diffusion type quantum Monte Carlo simulations. The main advantage of our strategy is that it reduces enormously computer memory. This, together with the surprisingly easy determination of the thresholds $W^\pm$, could make PERM also a very useful strategy for quantum Monte Carlo simulations.

In section 2 we will apply PERM to single homopolymers on the simple cubic lattice near their theta collapse temperature. This will be extended to semi-dilute solutions in Sec. 3, where we shall discuss critical unmixing of very long polymer chains. Finally, we shall sketch some further applications and improvements in Sec. 4. This includes what we call 'markovian anticipation' and applications to $2-d$ SAW's, to polymers in confined geometries, and to branched polymers.

2 The Θ Collapse of Single Chains

The Θ collapse is a tricritical phenomenon, described by the $n \to 0$ limit of the tricritical Landau-Ginzburg $O(n)$ model.[22] Its upper critical dimension is, as for all tricritical $O(n)$ models, $d = 3$. Therefore we must expect mean-field critical exponents in $d = 3$, modified only by logarithmic corrections. These corrections have been studied in great detail, partly due to some long-lasting controversies. The nowadays accepted results by Duplantier[23] include among others the following predictions:

(i) The average end-to-end distance for a chain of N elements grows at $T = T_\theta$ as

$$R_N^2/N \approx const \left(1 - \frac{37}{363 \ln N}\right) \tag{6}$$

(ii) The specific heat per monomer scales, again exactly at $T = T_\theta$, as

$$c(T = T_\theta) = \frac{1}{N}\left(\frac{1}{k_B T_\theta}\right)^2 (\langle m^2\rangle - \langle m\rangle^2) \sim (\ln N)^{3/11} \tag{7}$$

where m is the number of non-bonded nearest neighbor contacts.

(iii) In the collapsed phase $T < T_\theta$, the chain forms a globule with a bulk monomer density ϕ which scales as

$$\phi \sim (T_\Theta - T)[-\log(T_\Theta - T)]^{7/11}. \tag{8}$$

Here we shall only discuss chains modeled as SAW's on the simple cubic lattice with nearest neighbor attraction. For simplicity we choose this energy as $\epsilon = -1$, so that each non-bonded next neighbor contact contributes a factor $q = e^\beta$.

Two-dimensional chains are treated with similar algorithms in Ref.[19,20], chains on other 3-d lattices and off-lattice chains are discussed in Ref.[9,21].

The simulations presented below show most dramatically the power of PERM. This depends mainly on the fact that the above logarithmic corrections are weak, and thus a simple random walk needs very little pruning and enrichment. To a very good approximation, the chain length n performs during the simulation a random walk. Each addition of a monomer is a forward step of this walk, each pruning event is a jump backward to the next value of n on the stack. The efficiency of the algorithm is directly measured by the diffusion constant D of this walk, defined by

$$\langle (n_2 - n_1)^2\rangle = 2Dt, \tag{9}$$

where t is the number of steps (=subroutine function calls in the recursive implementation of Ref.[9]) performed between lengths n_1 and n_2. The larger D, the faster a long chain is built up and disassembled again. Measurements gave $D \approx 2000$ right at $T = T_\theta$. This is to be compared to $D = O(1)$ for other chain growth algorithms.[24,25]

Results of simulations with $N = 10,000$ are shown in Figs. 1 and 2. Both demonstrate clearly that the dominant behavior is mean field like, but the detailed predictions of Eqs.(6) and (7) are not verified. Indeed, the corrections are much

larger than predicted. It is not clear what is the source of these problems. It could be that higher order logarithmic corrections are important. But it could also be that multi-body forces between more than 3 monomers are important. Asymptotically, for $N \to \infty$, one expects three-body forces to be the only relevant ones (and this underlies Eqs.(6) and (7)), but at finite N this might not be true. In any case, any

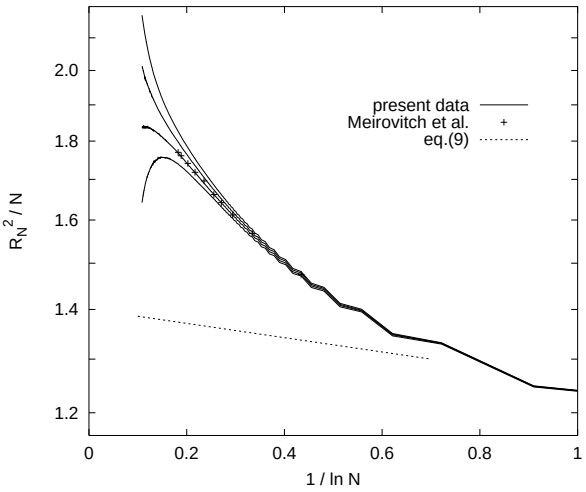

Figure 1: End-to-end swelling factor R_N^2/N against $1/\ln N$, compared to previous[27] MC estimates ($\diamond$) and to leading logarithmic prediction (dotted line). Curves are for $e^{1/kT} = 1.300$, 1.305, 1.310, 1.315 (top to bottom).

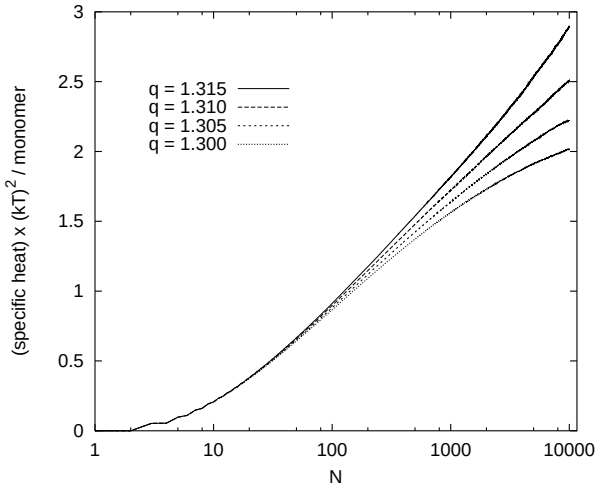

Figure 2: Variance of contact energies ($\sim$ specific heat) per monomer, for the same temperatures as in Fig. 1 ($q = e^{1/kT}$). Theory predicts $C/N \sim (\log N)^{3/11}$.

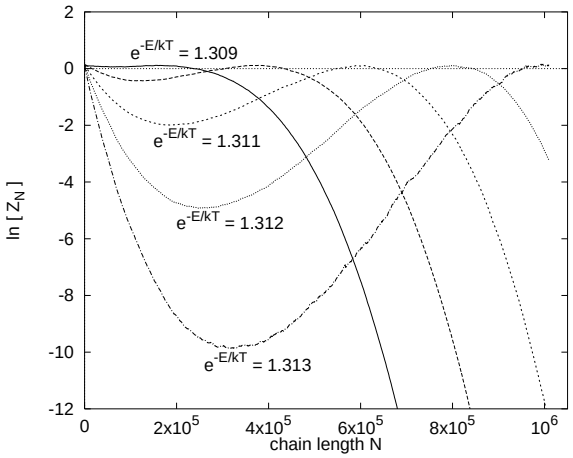

Figure 3: Logarithms of partition sums (negative free energies) of single chains in a box of size $256 \times 256 \times 256$. Notice that these are total free energies, not free energies per monomer. Fugacities are adjusted such that both maxima have the same height. The locations of the right-hand maxima give then estimates of $\Phi_\infty(T)$.
Results of this procedure are shown in Fig. 4.

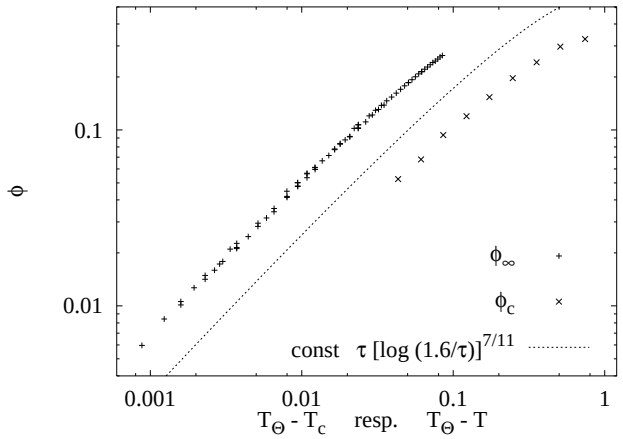

Figure 4: Plus signs (+): log-log plot of the infinite-N monomer density ϕ against $T_\theta - T$. Crosses ($\times$): log-log plot of critical monomer densities at finite N, plotted against $T_\theta - T_c(N)$, where $T_c(N)$ is the $N-$dependent critical temperature for unmixing.
Dashed line: theoretical prediction, up to arbitrary factor.

attempt[26] to interpret present experimental data by means of Eqs.(6) and (7) is likely to give wrong conclusions.

In order to test prediction (iii) directly, one would have to simulate extremely long chains ($N \gg 10^6$), since otherwise results are strongly influenced by the surface of the globule. As an alternative, one can simulate systems without any surface by squeezing a chain into a cube with periodic boundary conditions. Above the Θ-

point, the free energy will be a (cap-)convex function of N, while it should be non-convex for $T < T_\theta$. From this it should be possible to estimate both T_θ and the density ϕ. We were able to simulate chains of length up to 1.6×10^6, in lattices of up to 2^{25} sites. Some typical results are shown in Fig. 3. They show that T_θ can indeed be estimated reliably, with results agreeing perfectly with those from Figs. 1 and 2. Values of ϕ obtained from Fig. 3 and similar plots are shown in Fig. 4, together with critical densities of multichain systems discussed in the next section. The dotted line in Fig. 4 represents the prediction (8). We see very good agreement, much better than expected after having seen Figs. 1 and 2.

3 Critical Unmixing

Far below the Θ collapse temperature, it is intuitively obvious that not only different parts of one long chain will coalesce, but also different chains will lump together and precipitate. For finite but large N one expects therefore an unmixing transition at a critical temperature $T_c(N)$ which approaches T_θ from below when $N \to \infty$. For any fixed N this will be a critical point in the Ising universality class, but additional scaling laws are expected as $N \to \infty$. A schematic phase diagram is shown in Fig. 5.

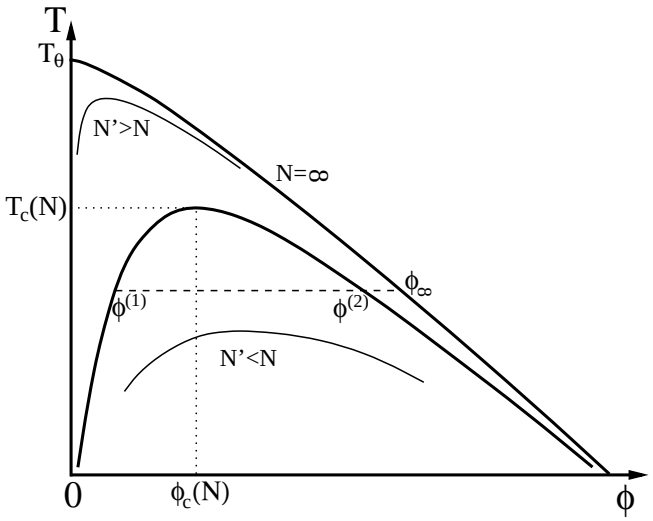

Figure 5: Schematic phase diagram for semi-dilute solutions of chain polymers. Uppermost curve: monomer concentration inside an infinitely large collapsed globule under zero outside pressure. Lower curves: coexistence curves for fixed chain length N. Curves are strictly monotonically ordered, with the coexistence curve for N below that for N' if $N < N'$.

For large N one expects critical densities $\phi_c(N)$, deviations of $T_c(N)$ from T_θ, and critical amplitudes $\phi^{(2)} - \phi^{(1)}$ to scale with N,

$$\phi^{(2)} - \phi^{(1)} \sim (T_c - T)^\beta N^{-x_1} \,, \tag{10}$$

$$\phi_c \sim N^{-x_2} \,, \tag{11}$$

$$T_\theta - T_c(N) \sim N^{-x_3} \ . \tag{12}$$

Flory-Huggins theory[28] predicts $x_1 = 1/4,\ x_2 = x_3 = 1/2$. Notice that these predictions stand on rather different footings. Exponents x_2 and x_3 actually describe only properties which are not influenced by critical Ising-type fluctuations, and should therefore be correctly predicted by Flory-Huggins theory. In contrast, one has to expect that x_1 is heavily influenced by anomalous Ising scaling, since it appears in conjunction with β which is different from its mean-field value in $d = 3$. Indeed we had problems in measuring x_1 reliably.[29] Therefore we shall discuss in the following only x_2 and x_3, and the chain swelling R_N^2/N. Naively one should expect that chains are ideal at the critical point.

Experiments[30,31,32,33,34,35] suggest very strongly that x_3 is smaller than its mean-field value, $x_3 = 0.38 \pm 0.01$. This problem has lead to numerous speculations.[36,38,37] In particular, it was suggested [36,37,38,39] that chains are already partly collapsed at $T_c(N)$. For a review, see Ref.[40].

Simulations with PERM are straightforward. We just have to place every N-th monomer not next to the previous one but at a random free site. In order to get the correct statistics (different chains are identical) we have to divide the partition sum by a factor $1/M!$ for a system with M chains. We simulated chain lengths between $N = 8$ and $N = 4096$. The latter is much larger than in any other simulation (the longest being $N = 1000$ in Ref.[41]).

Except for $N = 4096$, lattice sizes were such that the critical systems had roughly 50 to 100 chains. Although this is at least as large as in most previous simulations, it is not enough to be free from very large finite size corrections. In order to cope with them, we used a double histogram method as suggested in Ref.[42] and applied to the present problem in Ref.[43,41]. The main advantage of this method is that it allows to use field mixing[42] to get rid of the very strong asymmetry of the histogram and to reduce the problem essentially to the well understood Ising problem where the symmetry $m \to -m$ allows to pin down very precisely the critical point. But the efficiency of field mixing was checked in Ref.[43,41] only in so far as it rendered the distribution of chain numbers M symmetric. The stronger requirement that also the joint (M, E)-distribution (E=energy) should be symmetric was not tested in Ref.[43,41] due to lack of sufficient statistics. In our simulations[29] we found that this symmetry is indeed badly violated. We have no explanation for this, and no remedy. It is mostly this problem which prevented us from obtaining reliable estimates of x_1, see the discussion in Ref.[29]. But other findings are much less affected by it.

Our first result is that chains are not collapsed at the critical point but slightly swollen. This agrees with Ref.[43]. The swelling is only weakly dependent on the monomer concentration ϕ. Results obtained exactly at the critical point are shown in Fig. 6. They show that the swelling agrees in the limit $N \to \infty$ perfectly with that of isolated chains.

Log-log plots of $\phi_c(N)$ versus $T_\theta - T_c(N)$, of $(T_\theta - T_c(N))/T_c(N)$ versus N, and of $\phi_c(N)$ versus N are shown in Figs. 4, 7, and 8. From Fig. 7 we see clearly that $x_3 \approx 1/2$, with surprisingly small logarithmic corrections. This might be due to our choice of scaling variable. Using $(T_\theta - T_c(N))/T_\theta$ instead of $(T_\theta - T_c(N))/T_c(N)$, though being equivalent for large N, would give much worse scaling. Nevertheless,

Fig. 7 suggests strongly that x_3 is equal to its mean field value. This is in marked contrast to x_2.

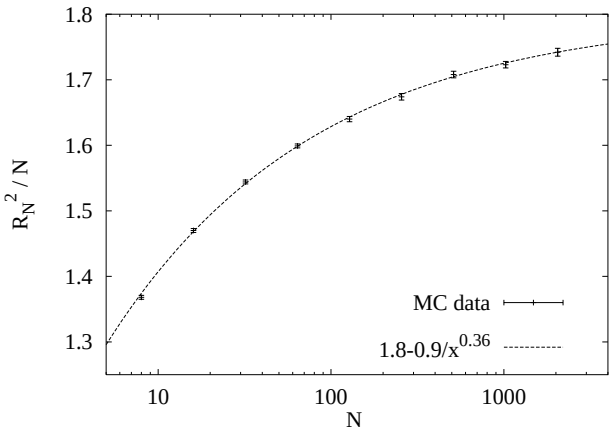

Figure 6: Swelling factors R_N^2/N at the critical point plotted against N. The dashed line is a fit with the function $1.8 - 0.9x^{-0.36}$. This fit has no particular significance except for the fact that the limit $N \to \infty$ agrees with the swelling of infinitely long free Θ polymers.

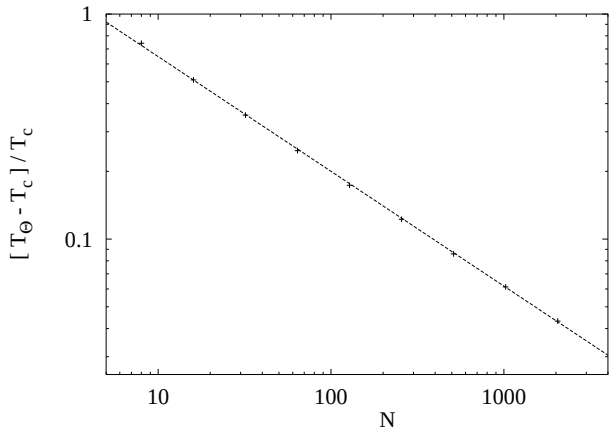

Figure 7: Log-log plot of $(T_\Theta - T_c(N))/T_c(N)$ against N (+). The dashed line has slope -0.51.

At first glance, Fig. 8 seems to give $x_2 = 0.385$, in perfect agreement with previous simulations[43,41] and with experiment.[30,31,32,33,34,35] But we see already in Fig. 8 that there are strong deviations at small values of N, suggesting that this is strongly affected by logarithmic corrections. This suspicion is confirmed by Fig. 4 and by a strong theoretical argument[29]: if $x_3 = 1/2$ (as confirmed by our simulations and by previous works[43,41]), and if coexistence curves are monotonically ordered as indicated in Fig. 5 (and as assumed by all previous authors), then x_2 is related to the dependence of $\phi(T)$ for isolated infinitely long chains: $\phi(T) \sim T^z$ with $z \leq 2x_2$.

Thus $x_2 < 1/2$ would be inconsistent with the generally accepted value $z = 1$ which we had also been confirmed by our simulations presented in Sec. 2. We thus conclude, in striking contrast to most previous authors, that also x_2 has its mean field value, and that all observed deviations are (expected!) logarithmic corrections. Further arguments in favor of this interpretation, based on an explicit evaluation of the free energy, are given in Ref.[29].

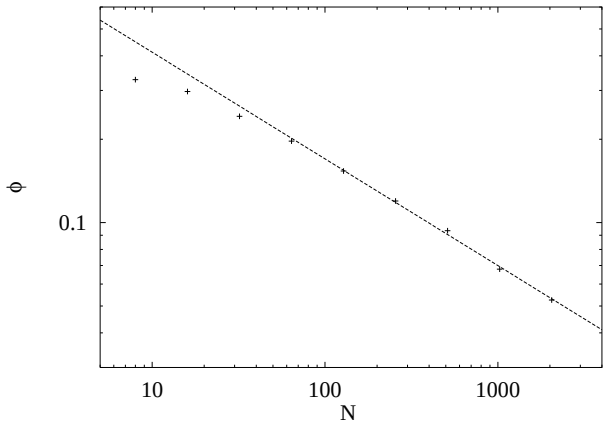

Figure 8: Log-log plot of $\phi_c(N)$ against N. The dashed line (slope -0.385) fits data for large N, but deviations at small N are substantial.

4 Further Developments

4.1 Markovian Anticipation and 2-Dimensional SAW's

As we had pointed out already in the introduction, PERM becomes most efficient if the a priori probabilities $p(\mathcal{C})$ are such that exactly the right distribution is obtained even without pruning and enrichment. In SAW's, this means that we should not choose randomly among the free neighbors when placing the next monomer. One way to "guide" the growth consists in looking ahead. In the 'scanning method' of Meirovitch[27] one looks k steps ahead by checking all possible k-step extensions, and decides on the success of these extensions which single step to take next. This is efficient, but also very time consuming: the effort increases exponentially with k.

An alternative which costs hardly anything in terms of CPU time (but which requires rather large memory, if pushed to the extreme) consists essentially in looking back k steps. Assume we are dealing with a lattice with coordination number $\mathcal{N}$. During the initial steps of the simulation we build up two histograms $H_0(i)$ and $H_1(i)$ of size $\mathcal{N}^{(k+1)}$ each. Here $i = 0, \ldots \mathcal{N}^{(k+1)} - 1$ labels all possible recent k-step histories and their one-step extensions. The first histogram $H_0(i)$ contains the weights of these $(k+1)$-step paths at the moment when the last step is made. The second histogram $H_1(i)$ contains the weights with which the same paths contribute much later (we used 100 steps in our simulations) to the partition sum. The ratio $H_1(i)/H_0(i)$ is therefore a direct estimate of how much the last step in path i con-

tributed to the final partition sum. Let us denote this last step as j $(= 0, \ldots \mathcal{N}-1)$ and the previous k steps as i_0, so that $i = (i_0, j)$. Then we propose

$$p(j|i_0) = \frac{H_1(i_0, j)/H_0(i_0, j)}{\sum_{j'} H_1(i_0, j')/H_0(i_0, j')} \tag{13}$$

as the best anticipation where to go next, based on a memory of k steps.

Equation gives indeed a substantial improvement for 2-d SAW's on all tested lattices (square, triangular, honeycomb, Manhatten). Using k such that $\mathcal{N}^{(k+1)} \approx 10^6$ we get in all cases an increase by roughly a factor 10 in the effective diffusion constant D (see Eq.(9)). This means that the number of independent chains is increased by roughly one order of magnitude over the uniform choice.

Applications and details will be published elsewhere.[44] Here we just mention that this allows e.g. very efficient simulations in confined geometries. To illustrate this, we show in Fig. 9 a single chain of length $N = 60,000$ in a slit of width 160. For convenience of plotting this was cut into 6 pieces. Such simulations allow significant tests of predictions about the monomer density in the vicinity of a hard wall.[45]

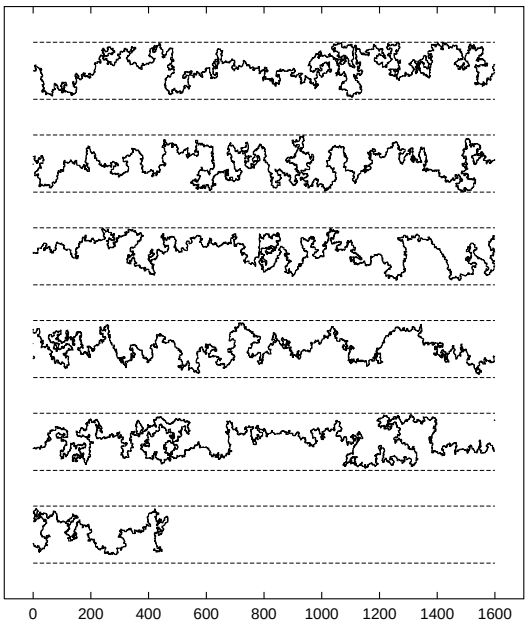

Figure 9: A single chain of length $N = 60,000$ in a strip of width 160 of a square lattice. In order to plot the configuration, it is cut in 6 pieces.

4.2 Lattice Animals

In our last example we want to illustrate that the strategy of PERM – starting off with a biased selection, weighting it properly, and using these weights to redirect the selection by cloning and killing – can also be implemented quite differently.

Consider the set of all connected clusters $\mathcal{C}_n$ of n sites on a regular lattice, with the origin being one of these sites, and with a weight $Q(\mathcal{C}_n)$ defined on each cluster. The (n-site) lattice animal problem is defined by giving the same weight to each cluster. The last requirement distinguishes animal statistics from statistics of percolation clusters. Take site percolation for definiteness, with 'wetting' probability p. Then a cluster of n sites with b boundary sites carries a weight $p^n(1-p)^b$ in the percolation ensemble, while its weight in the animals ensemble is independent of b. In the limit $p \to 0$ this difference disappears obviously, and the two statistics coincide. Due to universality, we expect indeed that the scaling behavior is the same for any value of p less than the critical percolation threshold p_c.

While there exists no simple and efficient algorithm for simulating large animals (for a recent Rosenbluth type algorithm see Ref.[46]), there exist very simple and efficient algorithms for percolation clusters. The best known is presumably the Leath algorithm[47] which constructs the cluster in a "breadth first" (i.e., layer by layer) way, the easiest to implement is a recursive "depth first" algorithm.[48]

Our PERM strategy consists now in starting off to generate subcritical percolation clusters by the Leath method, and in making clones of those growing clusters which contribute more than average to the animal ensemble.[49] Since we work at $p < p_c$, each cluster growth would stop sooner or later if there were no enrichment. Therefore we do not need explicit pruning. The threshold W^+ for cloning is chosen such that it depends both on the present animal weight and on the anticipated weight at the end of growth.

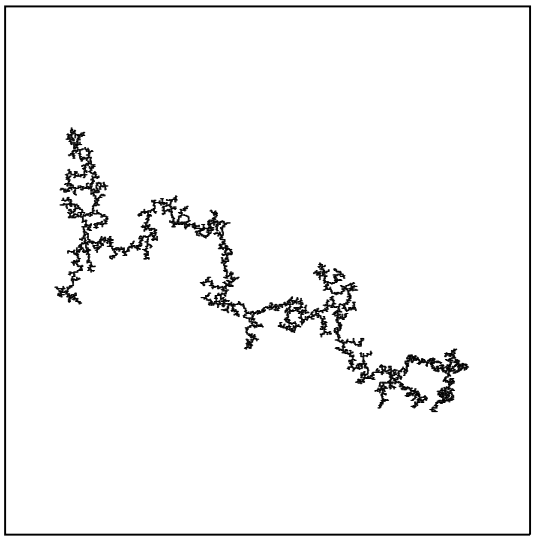

Fig. 10: A typical lattice animal with 8000 sites on the square lattice.

Usually, with growth algorithms like the Leath method, cluster statistics is updated only after clusters have stopped growing. But, as outlined below, one can also include contributions of still growing clusters. For percolation, this reduces slightly

the statistical fluctuations of the cluster size distributions, but the improvement is small. On the other hand, this improved strategy is crucial when using PERM to estimate animal statistics.

Consider a growing cluster during Leath growth. It contains n wetted sites, b boundary sites which are already known to be non-wetted, and g boundary sites at which the cluster can still grow since their status has not yet been decided ("growth sites"). This cluster will contribute to the percolation ensemble only if growth actually stops at all growth sites, i.e. with weight $(1-p)^g$. Since the relative weights of the percolation and animal ensembles differ by a factor $(1-p)^{(b+g)}$ (since now $b+g$ is the total number of boundary sites), this cluster has weight $W(\mathcal{C}) \propto (1-p)^{-b}$ in the animal ensemble. If we would use only this weight as a guide for cloning, we would clone if $W(\mathcal{C})$ is larger than some W^+ which is independent of b and g, and which depends on n in such a way that the sample size becomes independent of n. But clusters with many growth sites will of course have a bigger chance to keep growing and will contribute more to the precious statistics of very large clusters. It is not a priori clear what is the optimal choice for W^+ in view of this, but numerically we obtained best results for $W^+ \propto (1-p)^g$.

In this way we were able to obtain good statistics for animals of several thousand sites, independent of the dimension of the lattice. A typical 2-d animal with 8000 sites is shown in Fig. 10. We were also able to simulate animal collapse (when each nearest neighbor pair contributes $-\epsilon$ to the energy), and animals near an adsorbing surface. Details will be published in Ref.[49].

5 Discussion

We have presented a general strategy (PERM) for sampling from any given probability distribution. The main idea is to follow initially a biased distribution which is more easy to simulate. This bias is taken into account by re-weighting the sample, and these weights are used to interfere by pruning and cloning. For this to be possible it is needed that construction of instances is done in many small steps (which is always the case in statistical physics), and that the initial growth of weights is not misleading. It is mainly the second requirement which limits the usefulness of our approach. In the 3-d Ising model, for instance, our approach gave rather mediocre results. But in other cases it is much more efficient. For polymers near the Θ collapse point, in particular, it seems by far the most efficient method presently known.

The method shows some similarity to genetic and other evolutionary algorithms since good ('fit') instances are multiplied, while bad ('unfit') ones are killed. But in contrast to them this is done such that the wanted probability distribution is respected exactly, and it is done without the need for keeping large populations of instances in computer memory. We believe that it is mainly these features which make our strategy a promising candidate for further applications. Several such applications to toy protein models are discussed in these proceedings.[10,11,50] Further success will depend strongly on whether good non-trivial choices for $p(\mathcal{C})$ (such as in markovian anticipation, Sec. 3a) or for $W^{\pm}$ (such as in Sec. 3b) can be found in the specific problem at hand. Another promising avenue to follow is the combination

of our strategy with more conventional (Metropolis type) concepts. If this merge succeeds, PERM might lead to an efficient algorithms for finding native states of real proteins.

The authors are most grateful to Gerard Barkema for numerous discussions.

References

1. N. Metropolis *et al.*, J. Chem. Phys. **21**, 1087 (1953)
2. R.H. Swendsen and J-S. Wang, Phys. Rev. Lett. **58**, 86 (1987)
3. N. Madras and A.D. Sokal, J. Stat. Phys. **50**, 109 (1988)
4. B. Berg and T. Neuhaus, Phys. Lett. **B 285**, 391 (1991); Phys. Rev. Lett **69**, 9 (1992)
5. E. Marinari and G. Parisi, Europhys. Lett. **19**, 451 (1992)
6. M.N. Rosenbluth and A.W. Rosenbluth, J. Chem. Phys. **23**, 356 (1955)
7. W. Bruns, Macromolecules **17**, 2826 (1984)
8. F.T. Wall and J.J. Erpenbeck, J. Chem. Phys. **30**, 634, 637 (1959)
9. P. Grassberger, Phys. Rev. **E56**, 3682 (1997)
10. H. Frauenkron, U. Bastolla, E. Gerstner, P. Grassberger, and W. Nadler, these proceedings
11. U. Bastolla and P. Grassberger, these proceedings
12. J. Holland, *Adaptation in natural and artificial systems* (The University of Michigan, Ann Arbor, MI, 1992)
13. T. Garel and H. Orland, J. Phys. **A 23**, L621 (1990)
14. P.G. Higgs and H. Orland, J. Chem. **95**, 4506 (1991)
15. B. Velikson, T. Garel, J.-C. Niel, H. Orland, and J.C. Smith, J. Comput. Chem. **13**, 1216 (1992)
16. H. Orland, these proceedings
17. C.J. Umrigar, M.P. Nightingale, and K.J. Runge, J. Chem. Phys. **99**, 2865 (1993)
18. D. Aldous and U. Vazirani, *"Go with the winners" algorithms*; in Proc. 35th IEEE Sympos. on Foundations of Computer Science (1994)
19. P. Grassberger and R. Hegger, Journal de Physique **I 5**, 597 (1995)
20. G.T. Barkema, U. Bastolla, and P. Grassberger, cond-mat/9707312 (1997)
21. P. Grassberger and R. Hegger, Condens. Matter (J. Phys. C) **7**, 3089 (1995)
22. P.-G. de Gennes, J. Physique (Paris) Lett. **36**, L55 (1975); J. Physique (Paris) **39**, L299 (1979)
23. B. Duplantier, J. Chem. Phys. **86**, 4233 (1987); Europhys. Lett. **1**, 491 (1986); J. Physique (Paris) **43**, 991 (1982)
24. S. Redner and P.J. Reynolds, J. Phys. **A 14**, 2679 (1981)
25. A. Beretti and A.D. Sokal, J. Stat. Phys. **40**, 483 (1985)
26. A.T. Boothroyd, A.R. Rennie, C.B. Boothroyd and L.J. Fetters, Phys. Rev. Lett. **69**, 426 (1992)
27. H. Meirovitch and H.A. Lim, J. Chem. Phys. **92**, 5144 (1990)
28. P.G. de Gennes, *Scaling Concepts in Polymer Physics* (Cornell Univ. Press, Ithaca, 1979)

29. H. Frauenkron and P. Grassberger, J. Chem. Phys. **107**, 9599 (1997)
30. R. Perzynski, A. Delsanti, and M. Adam, J. Physique (Paris) **48**, 115 (1987)
31. T. Dobashi, M. Nakata, and M. Kaneko, J. Chem. Phys. **72**, 6685, 6692 (1980)
32. K. Shinozaki, T.V. Tan, Y. Saito, and T. Nose, Polymer **23**, 728 (1982)
33. B. Chu and Z. Wang, Macromolecules **21**, 2283 (1988)
34. K.-Q. Xia, C. Franck, and B. Widom, J. Chem. Phys. **97**, 1446 (1992)
35. Y. Izumi and Y. Miyake, J. Chem. Phys. **81**, 1501 (1984)
36. I.C. Sanchez, J. Appl. Phys. **58**, 2871 (1985)
37. M. Muthukumar, J. Chem. Phys. **85**, 4772 (1986)
38. I.C. Sanchez, J. Phys. Chem. **93**, 6983 (1989)
39. B.J. Cherayil, J. Chem. Phys. **95**, 2135 (1991); J. Chem. Phys. **98**, 9126 (1993)
40. B. Widom, Physica **A 194**, 532 (1993)
41. A.Z. Panagiotopoulos, V. Wong, and A. Floriano, cond-mat/9708095 (1997)
42. N.B. Wilding and A.D. Bruce, J. Phys. **C 4**, 3087 (1992)
43. N.B. Wilding, M. Müller, and K. Binder, J. Chem. Phys. **105**, 802 (1996)
44. P. Grassberger, to be published
45. E. Eisenriegler, *Random walks in polymer physics*; in "Field Theoretical Tools in Polymer and Particle Physics", H. Meyer-Ortmanns and A. Klümper, eds. (Springer, Heidelberg, 1997)
46. C.M. Care, Phys. Rev. **E 56**, 1181 (1997)
47. P.L. Leath, Phys. Rev. **B 14**, 5046 (1976)
48. J.-S. Wang and R.H. Swendsen, Physica **A 167**, 565 (1990)
49. P. Grassberger and W. Nadler, to be published
50. E. Gerstner, P. Grassberger, and W. Nadler, these proceedings